Collective Intelligence in Market-Based Social Decision Making

Research-in-Progress

Henner Gimpel **Florian Teschner**

Karlsruhe Institute of Technology (KIT)
Institute of Information Systems and Marketing
Englerstr. 14, 76131 Karlsruhe, Germany

henner.gimpel@kit.edu florian.teschner@kit.edu

Abstract

Public and private organizations increasingly engage in social decision making. Modern information systems facilitate participatory processes that harness the collective intelligence of crowds: Institutions open up for grassroots democracy and shift decisions towards a crowd of citizens, employees, or customers. Prediction markets are common tools in such social decision making. They have repeatedly proven their potency for aggregating information from a crowd of experts. However, their contemporary use in social decision making typically involves experts with a vested interest in the decision. They have incentives to manipulate the market and the decision. Simply carrying over prediction market knowledge to the design of social decision making might be premature. To back this hypothesis, we present findings from a lab experiment. Data suggests that people with stakes in the decision manipulate the market, corrupt information aggregation, and lessen decision quality. Common practice is better than random but worse than alternative designs.

Keywords: decision making, wisdom of the crowd, experimental economics

Introduction

Advances in information systems are changing decision making in many contexts: political institutions increasingly open up for grassroots democracy and open discussion of societal innovation, ad-hoc communities use social media to coordinate, and companies gradually shift decisions towards a broad basis of employees and allow for user-driven innovation. An underlying theme of this trend is using the collective intelligence and wisdom of the crowd. While the crowd is used to gather and aggregate information and make predictions about the future; the final decision typically remains with a principal e.g., the administration, city council, board of directors, or voters. A prominent example is the Policy Analysis Market (PAM): Starting 2001, the US Department of Defense planned an electronic market-based decision support system (Hanson 2006). PAM was, in essence, a prediction market to forecast policy-relevant events like riots and terror attacks. It failed due to ethical concerns with earning money by betting on disastrous events. A more recent and successful example is the US-based defense contractor Rite-Solutions: Any employee with an innovative idea can promote that idea in an idea market. The idea is presented via a company-internal initial public offering (IPO) and colleagues can trade virtual stocks in the idea (Lavoie 2009). The top ideas attracting most investments are subsequently sponsored by senior management. Upon success, the idea is taken from the market and all its investors are rewarded.

Prediction markets are not the only tool for social decision making, but a prominent one. The question is, however, how well they are suited for this task. On the positive side, they have proven repeatedly to be very potent information aggregation mechanisms (e.g. Berg et al. 2008; Ledyard et al. 2009; Teschner et al. 2011). On the other hand, participatory processes typically involve people with a vested interest in the decision. This induces manipulation of the market to influence the decision in one's own interest. An example is the Internet-based prediction market Intrade being temporarily manipulated in the 2012 US presidential election (Thomson 2012).

To understand the extent and influence of manipulation on information aggregation and decision quality, we study such markets experimentally. The contribution of our work is that we study informed manipulators with an interest in the decision taken subsequent to market-based information aggregation. In line with previous work (e.g. Deck et al. 2012), our experimental results show that manipulation impairs information aggregation. This study shows how manipulation affects the subsequent decision (quality). The extent of manipulation depends on the decision making rule – the current common practice performs better than random but worse than alternative decision making rules. These results add to the knowledge base of information systems research; designers can draw on these insights and apply them to future information systems for reshaping social decision making.

Background and Related Work

Prediction markets. Speculative markets aggregate information and deliver predictions on future events. Their predictions are oftentimes more precise than the outcome of other prediction mechanisms like polling or statistical models. Examples include sports events like soccer, weather forecasts, demand for natural resources, political elections, macro-economic indicators, value of new stocks to be issued as well as customer behavior and revenues. Ledyard et al. (2009) summarize: "In every known head-to-head field comparison, information markets have been no less accurate than other social institutions." Besides the academic interest in prediction markets, several organizations use prediction markets in their daily operations. This includes organizations like Chevron, Deutsche Bank, Eli Lilly, Ford, Goldman Sachs, Google, Hewlett-Packard, Intel, Microsoft, Procter & Gamble, and the RAND Corporation.

The key mechanism of a prediction market is as follows: A crowd is invited to trade virtual stocks, typically in an electronic market. The value of the stocks is tied to the outcome of a future event. Crowd members are paid according to their portfolio of stocks once the event is realized. Until then, prices are assumed to reflect the crowd's aggregate estimate on the future event – a likelihood of realization, expected product revenues etc. The examples listed above show that such market-based predictions generally perform very well. However, only rarely prediction is an end by itself.

Decision markets. Typically, a prediction is only generated as input for a decision. While most of prediction market research omits this, there are some exceptions. In the economic literature, these are, for example, Berg and Rietz (2003) studying prediction markets as decision support systems and Hanson (2006) analyzing decision markets for policy advice. In the computer science literature, examples include Otham and Sandholm (2010) and Chen et al. (2011), both are theoretically modeling decision markets. The key difference between prediction markets and decision markets is broadening the scope of analysis and explicitly taking the decision into account. This opens up two challenges:

1. Markets need to be conditional on the decision. This creates incentives for market manipulation in order to maximize gains from trade (Otham and Sandholm 2010; Chen et al. 2011).
2. Experts might not be decision-agnostic, but might hold innate preferences over the decision. If so, they have an incentive to manipulate the market in order to manipulate the decision. Such a manipulation might be costly within the scope of trading in the market, but might maximize the sum of gains from trade and the decision's consequences.

In practice, the link from the crowd's prediction to a decision is seldom explicit. For political elections and sports events – the early application areas of prediction markets – there is no direct link. An indirect link might, occur, however; e.g. via herding effects, strategic voting, social pressure. For prediction markets on corporate or public policy matters there clearly is a link. Setting up and running a prediction market is costly, as typically financial incentives are provided to successful traders. Thus, it is not plausible that the information is irrelevant to the body sponsoring the market, e.g. a city council. They will take the information into account when deciding. As such, prediction markets frequently are employed as an element in social decision making. They turn into decision markets.

An example: Assume a city is planning to restructure their public transportation – their streetcar tracks in city center are overcrowded, streetcars are frequently delayed, accidents are common. There are two options: either to build new overground streetcar tracks farther away from city center or a subway. The aim is to find a cost-efficient solution increasing quality of life. Both options have a chance to achieve that aim, but neither is certain to do so. City council might decide to run a prediction market to improve the forecast on, e.g., expected costs, handicaps during year-long construction, and quality of the resulting solution. The crowd of citizens can certainly help in estimating these effects. Their participation in the process will likely increase consent with city council's decision. However, some will clearly have an interest in manipulating the market (and thus the decision). The tunneler, for example, might prefer the subway; so might citizens living close to the construction area for new streetcar tracks. In such a scenario, how well will a prediction market reflect the collective best estimate on the effectiveness of the options?

Manipulation. Manipulation is influencing market prices to knowingly wrong levels. Allen and Gale (1992) considered three types of manipulation: action-based (changing the underlying fundamentals), information-based (spreading false information), and trade-based (buying, selling of shares). We are studying trade-based manipulation only. Deck et al. (2013, p. 49) provide a recent review on manipulating prediction markets. A common line of argumentation is that a manipulator is just another kind of noise trader (Hanson 2006; Hanson and Oprea 2007). If well-informed traders seized the opportunity to profit from manipulative trades, market accuracy might actually increase. If manipulation occurs, it is typically only short-lived (Wolfers and Zitzewitz 2004). The prevailing opinion is that long-term price manipulation does not occur or succeed and that manipulators are unable to distort price accuracy (Hanson et al. 2006; Berg et al. 2008; Oliven and Rietz 2004). Wolfers and Zitzewitz (2004) concluded that there "have been several known attempts at manipulation of these markets, but none of them had much of a discernible effect on prices, except during a short transition phase". Deck et al. (2013, p. 48) summarize that "research suggests prediction markets are robust to manipulation attacks".

In two closely related experiments, Oprea et al. (2007) and Deck et al. (2013) incentivize traders to manipulate prediction market prices. They have other subjects in the role of forecasters. Forecasters do not trade in the market but observe market prices and are asked to forecast the true value of the underlying asset. The core question of both these experiments is whether manipulators manage to mislead forecasters. Oprea et al. (2007) find no evidence for successfully manipulation of prices. Hence, forecasters are equally accurate whether or not manipulators are present in the market. Deck et al. (2013) adapted the design by Oprea et al. (2007) in a way as to make manipulation of prices easier. The design changes include allowing for short selling, perfectly informing manipulators on the asset's true value, and incentivizing manipulators only on wrong forecasts, not at all on gains from trade. These design changes proved successful: Deck et al. (2013) report effective price manipulation and misleading of forecasters. Besides Hansen et al. (2004), this is one of the very rare examples of successful manipulation of prediction markets.

Our experiment differs from the work by Oprea et al. (2007) and Deck et al. (2013) in a couple of important ways: First, we study conditional market, i.e. two parallel markets of which only one will be pay-off relevant in the end, depending on the principal's decision. This creates additional incentives for market manipulation (Otham and Sandholm 2010; Chen et al. 2011) that are not present in the work by Oprea et al. (2007) and Deck et al. (2013). Second, in their experiments manipulators are explicitly assigned to their role and are provided with a direction for price manipulation that is ex-ante uncorrelated with the true value of the asset. Our subjects decide themselves whether to manipulate and, if so, in which direction to manipulate. Third, we study thinner markets with two rather than eight traders. Fourth, we employ a different market mechanism, i.e. a logarithmic market scoring rule rather than a double auction with open order book. Fifth, the role of forecasters taken by subjects in the experiments by Oprea et al. (2007) and Deck et al. (2013) compares to the principal in our experiment, which is automated with decision rules that are common knowledge for subjects. In summary, this setting is distinctly different from prior work. We observe successful manipulation in some of our treatments under specific conditions. This adds to the short list of reports on successful manipulation of prediction markets. Overall, our results contribute to sharpening the understanding on when manipulation of prediction markets is successful and how negative the result of such manipulation is for decision makers.

Experiment Design

The experiment compares information aggregation and decision quality in an abstract induced preferences setting: There are two binary lotteries represented by two urns, A and B, holding 10 balls each. Per urn, 2, 5, or 8 of these are black, the others white. A principal – who will draw one ball from one of the urns – is interested in drawing a black ball. He decides which urn to draw from, the draw itself is random. The principal has a-priori knowledge on the potential states of nature (2, 5, or 8 of 10 balls per urn being black) but does not know the exact number of black balls in either of the urns. To gain further information prior to deciding, he runs two parallel prediction markets for experts to share their private information. One market for urn A, one for B. The market price is assumed to reflect the aggregate prediction of the probability to draw a black ball from the respective urn. Experts are financially compensated based on their trading performance. Recall the streetcar/subway scenario from the previous section: the two urns represent the two options, drawing a black ball stands for the desired outcome to have a cost-efficient improvement of quality of life. City council is the principal (on behalf of all citizens); citizens participating in the market are the 'experts'. One might question whether in real-world applications the potential true states of the world are fully known to the decision maker, as they are in our experimental setting. As assumed in most models and practical guidelines for decisions under certainty or risk, we assume that a fair representation of all states of the world can be achieved by the decision maker during his decision process. Narrowing it down to 9 states of the world (all combinations of 3 states for each of 2 urns) is a simplification to limit the complexity of the experimental task.

In the experiment, human subjects take the role of experts. In a partner matching (like Healy et al. 2010), a cohort of two subjects (like Jian and Sami 2012) interacts in a combination of parallel markets A and B. Each cohort is independent of any other cohort to increase statistical power. Subjects become 'experts' via private information (like in the experiments by Oprea et al. 2007, Healy et al. 2010, Jian and Sami 2012, Deck et al. 2013 and many others before that). Once the number of black balls in an urn is randomly determined, each subject receives a private signal that one of three a-priori options is not true. Subjects' private signals differ, i.e. in aggregate the group has certainty on the number of black balls in the urn. The same holds true for the other urn. The method of providing private information is analogous to the approach by Plott and Sunder (1988) in an influential information aggregation experiment, so is the information structure of market participants jointly holding all information to figure out the true state of nature.

The number of black balls in urns A and B is unrelated to each other, so are the private signals. Subjects have full knowledge of the rules of the game including this information structure. An example: 2 black balls in urn A, 5 in B. Subject 1 might get the information: "In urn A there are NOT 8 black balls; it must be 2 or 5. In urn B there are NOT 2 black balls; it must be 5 or 8". Subject 2 would receive the analogous private signal "A … NOT 5 … 2 or 8. B … NOT 8… 2 or 5". Experts can only interact via the prediction market.

Contrary to the forecasters in the experiments by Oprea et al. (2007) and Deck et al. (2013), in our experiment the principal is automated and his decision rule depends on the experimental treatment. We compare three treatments – termed A-PRIORI, PROBABILISTIC, and DETERMINISTIC – in a between subject design. Subjects participate in a series of 11 periods, one trial and 10 payoff-relevant periods. Each period follows 3 phases

1. *Private information and estimate:* The number of black balls per urn is determined and subjects receive private information. Each subject is asked for her private estimate on the probability of drawing a black ball from urn A and from urn B. Additionally each subject indicates which urn she would prefer. Truthful revelation is incentivized with a proper scoring rule (cf. Hanson 2003, p. 109). Private estimates are neither communicated to the other expert nor the principal.
2. *Prediction market:* The market uses a logarithmic market scoring rule (Hanson 2003; Jian and Sami 2012). Subjects can buy and sell virtual stocks for three minutes in two parallel markets. The

value of the stocks is linked to the color of the ball that will be drawn. The final market price is used as the group's best predictor for the number of balls in the respective urn. Bennouri et al. (2011), Jian and Sami (2012), and others use similar approaches. Theoretically, a risk neutral trader who trades only once or is myopic maximizes expected utility by truthfully trading based on the private belief (Hanson 2003). For non-myopic traders, utility-maximizing strategies depend on the nature of private information (Chen et al. 2010) – while it is known that bluffing and delaying can be beneficial for traders, a complete characterization of equilibria is unknown (Jian and Sami 2012). Hence, it is necessary to empirically establish the benchmark of manipulation-free behavior (see A-PRIORI treatment in the following). Iyer et al. (2010) and Ostrovsky (2012) have shown that under a set of relatively mild conditions, equilibrium prices converge towards the correct posterior price.

3. *Decision:* The principal decides which urn to draw from. In the A-PRIORI treatment, he relies only on his a-priori knowledge on the likelihood to draw a black ball from either urn. By design, this is 50% for both urns. Thus, he chooses either urn with 50% probability. In the DETERMINISTIC treatment, he chooses the urn where the markets predict the higher likelihood of drawing a black ball, i.e. the urn with the higher final market price. In case of tied prices, the choice is based on a-priori knowledge. In the PROBABILISTIC treatment, the principal uses a logit decision function: each choice has positive probability but he is more likely to choose the urn with the higher market price; the likelihood increases with the price difference. Note that these decisions are the same irrespective of whether the principal is assumed to be risk neutral or risk averse. Markets are conditional on the principal's choice (Berg and Rietz 2003): For the urn chosen by the principal, a ball is drawn and its color determines the experts' compensation for trading. The other market is void.

   PROBABILISTIC mimics the current common use of prediction markets in practice. The markets' information is taken into account but the principal has discretion to decide otherwise. DETERMINISTIC is an alternative where the principal pre-commits to following the markets' suggestion. A-PRIORI serves as benchmark.

The interesting twist is that experts have preferences over the decision. They are not decision-agnostic. For each expert, the color of the ball drawn determines the outcome of a binary lottery. If the ball is drawn from urn A, the lottery has higher expected value and higher risk than if the ball is drawn from urn B. Thus, depending on an expert's risk aversion and assumed probabilities of drawing a black ball from either urn, she will prefer either a choice of urn A or B. This preference is privately elicited after providing private information. The lottery payoff is on average about 5 times higher than from the market. Thus, in the PROBABILISTIC and DETERMINISTIC treatments, experts have an incentive to manipulate the market price, forgo gains from trade, but more than compensate for this by playing their preferred lottery. In the A-PRIORI treatment, prices are unrelated to the decision – there is no incentive for manipulation. Subjects' behaviors in the A-PRIORI treatment serve as empirical benchmark for manipulation-free behavior.

Different behaviors can be termed "manipulation." Here we are only interested in a special case, namely in trade-based external manipulation, i.e. experts' strategic influence on the market price to influence the principal's decision (cf. Hanson et al. 2006). Internal manipulation aimed at profiting within the same market or set of markets – e.g., bluffing or delaying (Chen et al. 2010; Jian and Sami 2010) or manipulation induced by conditional markets (Chen et al. 2011) – is out of scope and constant across treatments.

162 undergraduate and graduate students were recruited as subjects. Each participated in exactly one cohort in exactly one treatment. Thus, we observed 27 independent cohorts per treatment, each trading in 10 periods (excluding the trial) of two parallel markets (A and B). Overall, this totals to 1,620 prediction markets, 540 per treatment. The experiment was run with a custom-made web application in a lab where subjects could neither see nor talk to each other during the experiment. Sessions lasted around 75 minutes. Payments were linked to individual performance in the experiment, the average payment was 17 USD (13 Euro), Standard Deviation 2.22 USD (1.70 EUR). Instructions for participants are available upon request.

Experiment Results

The first key question is whether markets aggregated information despite incentives for manipulation. The information contained in a forecast can be assessed by regressing actual values on predicted values (Fair-Shiller regressions: Fair and Shiller 1989). In our experiment, actual values are the true probabilities of drawing a black ball. One predictor for this probability is the market price. A second predictor is the average of experts' private estimates. Table 1 displays the result of three Fair-Shiller regressions, one per treatment: The market price contains information on the true probability ($\beta 1$ is significantly different from zero) for each treatment.[1] Thus, information aggregation took place in each treatment.

Result 1: Prediction markets aggregate information whether or not market participants have an interest in the subsequent decision.

Table 1. Measuring information contained in forecasts

Fair-Shiller regression estimates. Dependent variable: true probability

(Significance codes: '***' 0.001 '**' 0.01 '*' 0.05)

	(1) A-PRIORI	(2) PROBABILISTIC	(3) DETERMINISTIC
α Intercept	-0.350 ***	-0.164 ***	-0.207 ***
$\beta 1$ Market price	0.854 ***	0.546 ***	0.636 ***
$\beta 2$ Average estimate	0.752 ***	0.724 ***	0.713 ***
n[2]	540	540	540
Adjusted R^2	0.401	0.277	0.302

In the A-PRIORI treatment, without an incentive for manipulation, the market price is slightly more informative than the private estimates ($\beta 1 > \beta 2$ for model 1; in a bootstrapping test, the difference between the relative contributions of these regressors to the model's R^2 is not significant at the 10% level). In the PROBABILISTIC treatment, the market price is significantly less informative than the average estimate ($\beta 1 < \beta 2$ for model 2, significant at the 10% level). In the DETERMINISTIC treatment, finally, both forecasts are statistically equally informative (10% level). This is a first indication for information aggregation being corrupted when experts have an interest in the subsequent decision, at least in the PROBABILISTIC treatment.

Further regression analysis (data not shown) indicates that average estimates are equally good predictors in all treatments (as they should be, due to same incentives) while market price is a significantly better predictor in A-PRIORI than in either DETERMINISTIC or PROBABILISTIC (5% level). These findings are backed by the analysis of absolute and squared prediction errors, i.e. differences between true probability and average estimate as well as between true probability and market price (data not shown).

Result 2: Market manipulation with the aim to influence subsequent decisions partially corrupts information aggregation.

The second key question is how information aggregation affects decision quality. Table 2 lists the average probability of achieving the 'best outcome' by treatment. As 'best outcome', we define the outcome the principal would choose would he have full information on the true state of nature. It is only defined when

---

[1] The data set is panel data with multiple cohorts each playing 10 periods. For each regression model we tested for the existence of individual, time, or mixed fixed (F Test) or random (Breusch-Pagan Lagrange Multiplier Test) effects. The analysis revealed no such effects, thus pooling data is adequate.

[2] If both subjects submitted an estimate, the average is calculated. If only one subject provided an estimate, this is taken as aggregate estimate. If neither subject provided an estimate, the observation is dropped from the respective analysis. Differences to 540 observations per treatment for missing data, i.e. neither subject submitting an estimate. Alternatives would be to replace missing estimates by the a-priori estimate (50%) or discard these cases from all analyses. Both alternatives lead to the same qualitative results in all statistics.

the principal has a clear preference, not when he is indifferent. For A-PRIORI, the probability is 50% by design. PROBABILISTIC achieves a higher hit rate (61%), DETERMINISTIC an even higher rate (72%). These hit rates are pairwise significantly different (Fisher's exact test on 2x2 contingency tables, 5% level).

Interestingly, the hit rate depends strongly on preference alignment: when both experts' private preferences – would they have to choose a lottery for themselves – are aligned and in addition the principal would want this lottery (which he does not know a-priori but we analyze ex-post), the hit rate goes up to 88% in the DETERMINISTIC treatment, presumably as experts manipulate in the direction of the best outcome. When experts' preferences are split (one prefers A, the other B), the hit rate decreases, presumably as the experts' manipulation attempts cancel out. In case the experts' preferences are aligned but different from the principal's preferences (i.e. the best outcome), the hit rate in deciding based on market prices goes down to 50% – the principal could as well toss a coin rather than run a market. Presumably experts reinforce each other in manipulation and render the market price fully uninformative.

Table 2. Quantifying decision quality by treatment

Average probability of achieving the best outcome (number of observations)

Preference alignment	A-PRIORI	PROBABILISTIC	DETERMINISTIC
Overall	50% (180)	61% (180)	72% (180)
Experts and principal aligned	50% (45)	74% (55)	88% (58)
Experts split	50% (123)	56% (120)	67% (103)
Principal deviates	50% (12)	50% (5)	50% (19)

Table 2 is descriptive; the corresponding inferential statistics are provided by a left- and right-censored Tobit models regressing the probability to achieve the desired outcome on the interaction of treatment and preference alignment (cf. Kleiber and Zeileis 2008, p. 141-144). 'Period' runs from 1 to 10 to account for potential learning effects. Results of the maximum likelihood estimation are provided in Table 3; all conclusions are robust to alternatively using an ordinary least squares (OLS) regression.

Model 1 is estimated on data from all 3 treatments with A-PRIORI as baseline. It suggests that with a PROBABILISTIC decision rule, the principal's hit rate increases beyond the a-priori probability only if by chance the preferences of all experts and the principal himself are aligned. However, in this scenario, the principal is even better off with a deterministic decision rule ($\beta 6 > \beta 3$ in model 1; significance of this difference is shown by $\beta 6$ being significantly different from zero in model 3; only data from PROBABILISTIC and DETERMINISTIC included). In addition when experts' preferences are split, a deterministic decision rule makes the principal better off than a decision on a-priori probabilities ($\beta 5$ in model 1) or a probabilistic decision based on market prices ($\beta 5$ in model 2).

Result 3: When the principal bases his decision deterministically on market prices, he is more likely to achieve the best outcome than when he either probabilistically processes market-based information or decides on a-priori probabilities – the only exception to this rule is when experts' preferences for a decision are aligned but opposed to the principal's preferences.

Result 4: When experts' preferences for a decision are aligned but opposed to the principal's preferences, none of the investigated market-based decision mechanisms allows the principal to make a better decision than simply using a-priori estimates.

Conclusion and Future Work

Governments, companies, and other organizations are gradually changing their interaction with their stakeholders: citizens, employees, partners etc. Modern information systems increasingly facilitate participatory processes that harness the collective intelligence of crowds. Common tools for doing so are prediction markets. A large body of knowledge exists on such markets: it typically relies on the assumption that either prediction is an end by itself (there is no subsequent decision) or that market participants are decision-agnostic.

Table 3. Determining effect of treatment and preference alignment on decision quality
Censored Tobit regression estimates. Dependent variable: probability to achieve the desired outcome
(Significance codes: '***' 0.001 '**' 0.01 '*' 0.05)

	(1) A-PRIORI baseline	(2) PROBABILISTIC and experts split baseline	(3) PROBABILISTIC and preferences aligned baseline
α Intercept	0.473 ***	0.503 ***	0.683 ***
β1 PROBABILISTIC and principal deviates	0.008	-0.044	-0.224
β2 PROBABILISTIC and experts split	0.056	–	-0.180 *
β3 PROBABILISTIC and all aligned	0.237 ***	0.180 *	–
β4 DETERMINISTIC and principal deviates	-0.001	-0.059	-0.239
β5 DETERMINISTIC and experts split	0.266 ***	0.278 ***	0.098
β6 DETERMINISTIC and all aligned	0.663 ***	0.788 ***	0.608 ***
β7 Period	0.004	0.009	0.010
n	540	360	360
Log(scale)	-1.056 ***	-0.653 ***	-0.653 ***
Log-likelihood	-294.1	-286.7	-286.7

In this paper we argue that contemporary use of prediction markets as a building block in social decision making typically involves market participants with a vested interest in the decision. They have an incentive to manipulate the market in order to affect the decision. Simply carrying over knowledge on prediction markets to the design of such decision markets might be premature. To back this hypothesis, we presented preliminary findings from a lab experiment. Data suggests that indeed, people who have a stake in the decision manipulate the market and partially corrupt information aggregation. The common practice of taking a prediction market into account while retaining the discretion to decide otherwise turns out to be beneficial in only one specific case: When all market participants have the same preference for what the decision should be and this additionally aligns with the principal's preference, then a prediction market is better than having no information aggregation. In any other case, stakeholders might simply forgo the effort of designing a market and trading, as prices are uninformative. Deterministically tying the decision to the market outcome turns out to more often produce the desired decision. Credibly implementing this policy will, however, have practical hurdles.

Our results add to the literature on manipulation of markets, specifically manipulation of prediction markets. The prevailing opinion in theoretical, empirical, and experimental research on the topic is that manipulators typically do not succeed in systematically manipulating prices. Besides Hansen et al. (2004) in the field and Deck et al. (2013) in the lab, we present one of the rare examples of successful systematic manipulation of prediction markets. This contributes to the understanding on when manipulation of prediction markets is successful and how negative the result of such manipulation is for decision makers or other observers of such markets.

The limitations of the present work are straightforward and include the following: First and foremost, we approached the topic experimentally. By exploring market manipulation and decision quality while systematically controlling and varying environmental and institutional factors in the lab we acquire insight into when and why markets might fail to deliver valuable information. On the other hand, results have to be interpreted carefully, as details of real-world environments and institutional factors will matter. Results from the lab can inform behavioral mechanism design but cannot be directly transferred to the field. Second, the experiment design has only two participants per market. Thus, preference alignment among them is binary: alignment or split. This is a stress test for the market mechanism in a thin market. Real-world applications would likely involve more traders and might render manipulation attempts less successful unless jointly carried out by a large share of traders. Third, the experiment is conducted in the lab with student subjects in an abstract, induced-preferences setting. Both the subject pool and the setting are standard in experimental economics and both have been criticized and discussed extensively. We believe that the abstract setting provides greater control over subjects' believes than a framed setting would do, thus allowing for more rigorous analysis and interpretation. Student subjects have been used in countless experiments delivering valuable insights – however, conducting field studies, e.g. in corporate prediction markets, would certainly be a valuable robustness check for our results.

The reason for these limitations is that in the larger research program we first prioritized control and rigor over relevance. Future work should gradually relax these limitations. Specifically, theoretical modeling should be expanded to complement the experimental approach and field experiments should check the robustness of findings in framed settings with other subjects. Within the narrower scope of experiments, it will be valuable to study bluffing and delaying strategies in markets with market scoring rules, to study other forms of manipulation, and to further vary environmental and institutional factors like the information structure, the market mechanism, and the number of traders in a market to further sketch out the preconditions for manipulation being effective or not. Finally, the practical implications for designing and running information systems for using collective intelligence in social decision making need to be examined further.

References

Allen, F. and Gale, D. "Stock-Price Manipulation," *The Review of Financial Studies* (5:3), 1992, 503–529.

Bennouri, M., Gimpel, H. and Robert, J. "Measuring the impact of information aggregation mechanisms: An experimental investigation," *Journal of Economic Behavior & Organization* (78:3), 2011, 302–318.

Berg, J. E., Nelson, F. D. and Rietz, T. A. "Prediction market accuracy in the long run," *International Journal of Forecasting* (24:2), 2008, 285–300.

Berg, J. E. and Rietz, T. A. "Prediction Markets as Decision Support Systems," *Information Systems Frontiers* (5:1), 2003, 79–93.

Chen, Y., Dimitrov, S., Sami, R., Reeves, D., Pennock, D., Hanson, R., Fortnow, L. and Gonen, R. "Gaming Prediction Markets: Equilibrium Strategies with a Market Maker," *Algorithmica* (58:4), 2010, 930–969.

Chen, Y., Kash, I., Ruberry, M. and Shnayder, V. "Decision markets with good incentives," *Internet and Network Economics*, Springer, 2011, 72–83.

Deck, C., Lin, S. and Porter, D. "Affecting policy by manipulating prediction markets: Experimental evidence," *Journal of Economic Behavior & Organization* (85), 2013, 48–62.

Fair, R. C. and Shiller, R. J. "The informational context of ex-ante forecasts," *The Review of Economics and Statistics* (71), 1989, 325–331.

Hansen, J., Schmidt, C. and Strobel, M. "Manipulation in Political Stock Markets – Predictions and Evidence," *Applied Economics Letters* (11), 2004, 459–463.

Hanson, R. "Combinatorial Information Market Design," *Information Systems Frontiers* (5:1), 2003, 107–119.

Hanson, R. "Decision Markets for Policy Advice", *in* Patashnik, E. and Gerber, A., ed., *Promoting the General Welfare: American Democracy and the Political Economy of Government Performance*, Brookings Institution Press, Washington D.C., 2006.

Hanson, R. and Oprea, R. "A Manipulator Can Aid Prediction Market Accuracy," *Economica* (76), 2007, 304-314.

Hanson, R., Oprea, R. and Porter, D. "Information Aggregation and Manipulation in an Experimental Market," *Journal of Economic Behavior and Organization* (60:4), 2006, 449–459.

Healy, P. J., Linardi, S., Lowery, J. R. and Ledyard, J. O. "Prediction Markets: Alternative Mechanisms for Complex Environments with Few Traders," *Management Science* (56:11), 2010, 1977–1996.

Iyer, K., Johari, R. and Moallemi, C. C. "Information Aggregation and Allocative Efficiency in Smooth Markets" *Proceedings of the 11th ACM Conference on Electronic Commerce*, Association for Computing Machinery, New York, 2010, 199–206.
Jian, L. and Sami, R. "Aggregation and manipulation in prediction markets: Effects of trading mechanism and information distribution," *Management Science* (58:1), 2012, 123–140.
Kleiber, C. and Zeileis, A. *Applied Econometrics with R*, Springer, 2008.
Lavoie, J. "The Innovation Engine at Rite-Solutions: Lessons from the CEO," *The Journal of Prediction Markets* (3), 2009, 1–11.
Ledyard, J., Hanson, R. and Ishikida, T. "An experimental test of combinatorial information markets," *Journal of Economic Behavior & Organization* (69:2), 2009, 182–189.
Oliven, K. and Rietz, T. A. "Suckers Are Born but Markets Are Made: Individual Rationality, Arbitrage, and Market Efficiency on an Electronic Futures Market," *Management Science* (50:3), 2004, 336–351.
Oprea, R., Porter, D., Hibbert, C., Hanson, R. and Dorina, T. "Can Manipulators Mislead Market Observers?" ESI Working Paper, Chapman University, 2007.
Ostrovsky, M. "Information Aggregation in Dynamic Markets With Strategic Traders," *Econometrica* (80:6), 2012, 2595–2647.
Othman, A. and Sandholm, T. "Decision rules and decision markets," *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1*, International Foundation for Autonomous Agents and Multiagent Systems, 2010, 625–632.
Plott, C. R. and Sunder, S. "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica* (56:5), 1988, 1085–1118.
Teschner, F., Mazarakis, A., Riordan, R. and Weinhardt, C. "Participation, Feedback & Incentives in a Competitive Forecasting Community," *Proceedings of the International Conference on Information Systems (ICIS)*, Shanghai, China, Paper 16, 2011, 1–14.
Thomason, D. "Should Presidential Campaigns Spend More Money Manipulating Intrade?" *The Atlantic*, October 23, 2012.
Wolfers, J. and Zitzewitz, E. "Prediction Markets," *Journal of Economic Perspectives* (18:2), 2004, 107–126.